Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh

23rd Annual Progress Report June 2002

Chapter 10: Optimal Control of Delta Salinity

Author: Eli Ateljevich

10 Optimal Control of Delta Salinity

[NOTE: In a project funded by DWR's Delta Modeling Section, Section member Eli Ateljevich completed a civil engineering doctoral dissertation on December 19, 2001. This chapter is a summary of that work; a complete copy of his dissertation is on file with the University of California, Berkeley.]

10.1 Introduction

Efficient control of salinity in the Delta is a subject that pervades the work of the Department's Modeling Support Branch. Water quality constraints are important in current State Water Project (SWP) operations, and even more so under the high-growth scenarios considered in water project planning. Because of this broad context in which salinity control plays a role, studies of salinity compliance have traditionally been incorporated into a larger modeling problem that goes something like this:

- □ Create a model of water project operations (CALSIM), and
- □ Include a module to model salinity intrusion and stipulate minimum inflows required to meet salinity standards (ANN, G-Model, etc). This constraint will be active during a fraction of the simulation period.

Two important simplifications are represented in this approach. The first is that the salinity module is, by computational necessity, an empirical surrogate for a physical model such as DSM2. The second is that the inflow requirement is calculated one time step at a time. The flow history up to the current time step is accepted as a *fait acompli*, and the calculations for the current step do not reckon the costs of putting the system into a costly position later. In optimization, this is known as a *greedy* treatment of time. In the present case, it is compounded by the use of large computational time steps.

Ateljevich (2001) tackles salinity control using a tack that is complementary to the traditional one: the SWP context is simplified, but the physics and the time component are treated in more detail. Under these assumptions, the *optimal water cost compliance problem* is an optimal control problem, the solution to which is the optimal schedule of upstream releases into the Delta and exports. The technical details of this problem make up the bulk of the dissertation.

10.2 Minimum Water Cost Compliance

The minimum water cost compliance problem is posed over a fixed time interval of several weeks or months. Salinity appears as an important operating constraint, and stochastic influences are neglected. The features of the problem are as follows:

- □ The cost function is the sum or cost over time of upstream releases.
- □ The volume of pumping over the optimization interval is fixed as a constraint
- □ Salinity constraints are imposed at one or more monitoring stations in the interior of the Delta.
- □ Delta dynamics are modeled using physical conservation equations of mass, momentum, and salinity transport.
- □ Physical bounds are imposed on the control variables to represent pumping capacity, minimum outflow requirements, and uncontrolled Delta inflows.

Variants of this formulation can be devised. For instance, pumping can be priced as a benefit rather than being fixed as a constraint. The salinity constraint may also be imposed on instantaneous or period-averaged salinity (the latter is the important case in the Delta). Finally, additional *regularization* components may be required to make the problem mathematically well posed.

The results of the dissertation indicate a high degree of control over salinity, particularly over daily-averaged values. The water cost under optimization is about 10% better than the cost obtained under very careful trial-and-error experiments and about 15-20% better than that of a good guess. The minimum water cost can also help to refine definitions of water cost and carriage water, which are usually calculated with rather gross assumptions about the time trajectory of pumping and inflow. Carriage water calculations are shown to be quite sensitive to the assumption that both pumping and inflow are constant over time.

10.3 Optimal Control Solutions

Optimal control solutions can be either hard to obtain or erratic. The features that go into the formulation of a well-posed optimal control problem are discussed at length in Ateljevich (2001). One issue deserves highlighting because it is critical to understanding what types of solutions can be obtained. That issue is *consistency* – what happens to the solution as the space and time steps are refined.

When expressed in continuous time, the optimal water cost compliance problem requires an "infinite" number of decisions – upstream inflow and pumping must be determined at an infinite number of times. This level of detail is not needed in practice, but water modelers are accustomed to discrete models that, under refinement, home in on a unique, physical continuous

solution. Constrained optimal control problems are not naturally so well behaved, even when a consistent model is used for the physics. Unless the problem is properly regularized, the prescribed control (if it exists) may exhibit infinite switches between the minimum and maximum control (*chattering*) or other types of singular or impractical behavior. Discrete solvers trying to home in on such a degenerate solution will tend to stall in ways that are time-step specific and very difficult to diagnose (often the diagnosis is "local minima" even when there is no evidence that any sort of local minimum is actually achieved or even exists). Besides being impractical, such wildly fluctuating controls violate the gradually varying flow assumptions under which the flow model is formulated.

10.3.1 Parameterization

At the other end of the spectrum, the problem can be parameterized crudely by taking one control variable per large time step (say several days or weeks), and using a representation that is piecewise constant or linear. This characterization involves a small number of decision variables. The outcome of optimization over such a restricted control space will be physically reasonable because the parameterization is too crude to represent the sort of chatter and pulses that occur when optimal control solutions go awry of common sense. On the other hand, one control per week in a system dominated by 25-hour and 14-day cycles may not be refined enough – important characteristic time scales of the problem are ignored or crudely represented. The catch is that the problem cannot be refined too much without losing the stability the crude parameterization imparts on the system.

The foregoing discussion allows for physically meaningful solutions to be developed in two ways. First, we can simply toss aside the notion of the "underlying truth" and accept a parameterized solution to the unregularized problem. Formal convergence is unlikely – the solution will always be time-step dependent. However, it may be possible to verify using a sensitivity analysis that the solution is not erratically dependent on time step in some neighborhood of interest – say, comparing four-day controls to two-day controls. In this case, the character of the *actuators* of the controls (pumps, reservoir gates, etc.), which operate with restricted complexity, may physically justify the use of a crude time step.

10.3.2 Regularization

Alternatively, we can add small regularization terms to the continuous problem, so that it does have a well-behaved solution. One common way to avoid discontinuities is to add a very small quadratic penalty on the time derivative of the control variables, which penalizes extremely sharp jumps and is insignificant the rest of the time. It is then possible to develop a solution algorithm that is consistent with this regularized problem and stop at whatever level of time detail seems appropriate. The regularization approach leads to similar water cost as the parameterized approach, but not to similar control sequences (usually the regularized solution has smaller extreme values).

Regularization of the problem requires extra development time and software designed for optimal control, but when applied it is much quicker to converge. This technique cannot be applied to DSM2 because some details in the implementation of DSM2-QUAL make it difficult

to calculate the derivative of the objective function with respect to the control variables at small time steps. Such calculations are usually made using variational (adjoint) techniques. Without this possibility, finite differences approximations must be used, perturbing the controls one at a time. This method is slow and inexact compared to variational methods. Parallel computing of DSM2 can be used to reduce the calculation time enough to make finite differences practicable (Finch and Kao, 1992).

10.3.3 Comparing Parameterization and Regularization

Because of these limitations, the differences between regularization and parameterization cannot be investigated directly for DSM2, although they can be illustrated on simplified problems. Figures 10.2 and 10.3 shows results for an optimal water cost control problem on a contrived 1-channel domain (Figure 10.1). The simplified problem has an upstream inflow, mid-stream pump, period average salinity constraint, ocean boundary, and fairly large dispersion. The numerical model used for the hydrodynamics is a four-point box scheme similar to that used in DSM2-HYDRO. The numerical scheme for salinity transport is the Flux Based Modified Method of Characteristics with a MUSCL limiter as described by Roache (1992). Derivatives were obtained using variational methods. A minimum flow of 200 cfs was enforced for inflow (0 cfs for pumping); the upper bounds were very high and did not become active during the optimization.

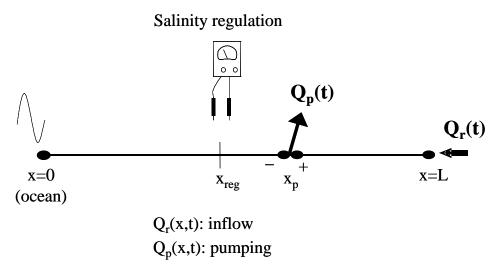


Figure 10.1: Sample 1-Channel Domain with Salinity Regulation.

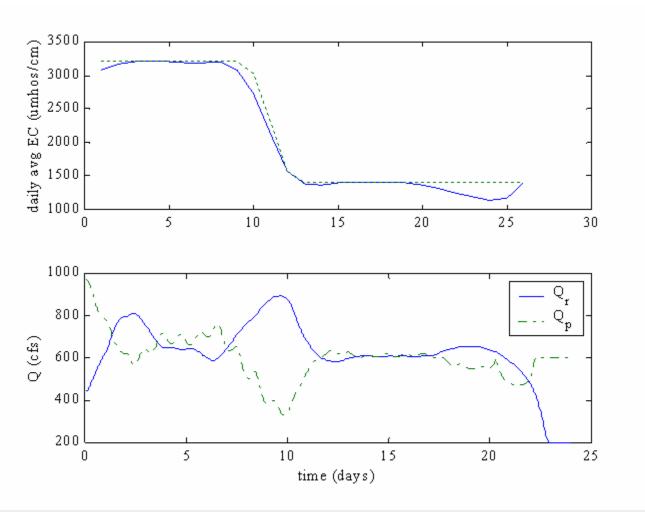


Figure 10.2: Regularization Solution to Meet Daily-Averaged Salinity Regulations with Control Trajectory for River Inflow (Q_r) and Pumping (Q_p) .

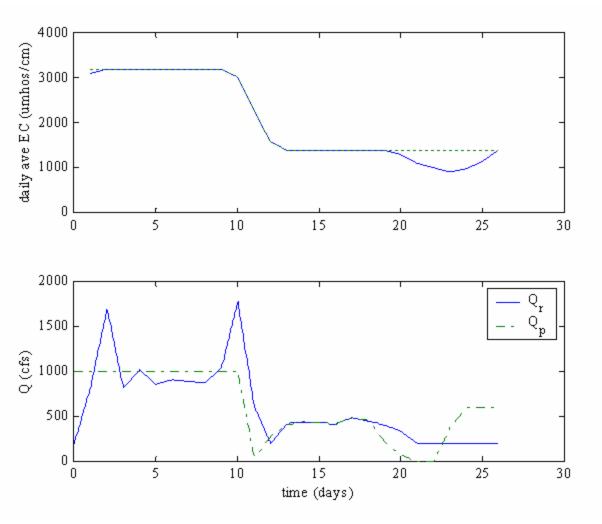


Figure 10.3: Parameterization Solution to Meet Daily-Averaged Salinity Regulations based on Piecewise Linear Controls for River Inflow (Q_r) and Pumping (Q_p) .

The control solutions in the parameterized and regularized cases are not similar in character – the parameterized solution reaches greater extremes than the regularized solution. However, the water cost of the two methods is extremely close, indicating that both methods are good in a bottom-line sense. The solution to the parameterized problem degrades to wild fluctuations at time steps of fewer than four hours (not shown); the solution to the regularized problem does not change significantly under refinement.

Two features visible in Figure 10.2 are typical and will be seen again in the DSM2 applications presented in Section 10.4. First, there is a very high degree of control on salinity. The piecewise linear daily controls are sufficient to allow the salinity solution to lie almost exactly on the regulation except during transitional periods (the beginning, end and times when the regulation is changing). Such "exact compliance" was not a goal of the optimization, but it is interesting to see that it can be achieved in a deterministic modeling setting.

Second, controls produced by the optimizer in the final few days of the problem have a *greedy* character. Water released in the final time steps will not affect salinity at the monitoring station until after the optimization is over. At the same time, this water has a cost. Under optimization, releases in the final time steps will therefore naturally be set to the minimum value. This often leads to a situation where salinity in the final time step is compliant, but increasing – an infraction in the next time step is imminent. Methods for dealing with end-of-period behavior are enumerated in some detail by Ateljevich (2001), and include constraints and penalties on the final value of salinity or its time derivative. These methods generally did not work as well as just padding the problem with some extra time. Fortunately, the earlier part of the example solution (before day 15) is not particularly sensitive to the way the end-of-period effects are handled.

10.4 Experiments Using DSM2

Only the parameterized control approach was available for optimization using DSM2, where finite differences estimate derivatives. Nevertheless, it is possible to obtain solutions that are reasonably stable and exhibit most of the behavior expected from experiments using a more rigorous approach on a simplified domain.

As an example problem, Ateljevich (2001) used parameterized controls to examine a 1994 compliance problem (Creel, 2000), in which reliance on an incorrect gage allowed salinity to approach the legal limit at the Contra Costa Water District pump intake at Rock Slough (Figure 10.4). The error was discovered and pumping was curtailed for a period of about one week to bring salinity levels back down. The Sacramento River inflow and combined export pumping rates are shown in Figure 10.5. The Delta Cross Channel was open the entire time. As suggested by Figure 10.4, the pulse-like reduction in pumping overcompensated, causing salinity to drop by half.

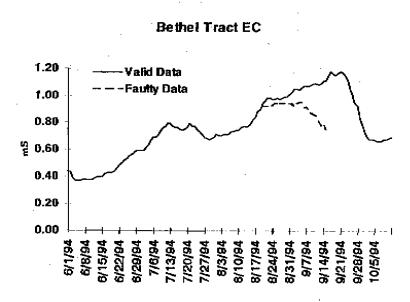


Figure 10.4: Bethel Tract Daily-Averaged EC (in mS) in Fall 1994. (taken from Creel, 2000)

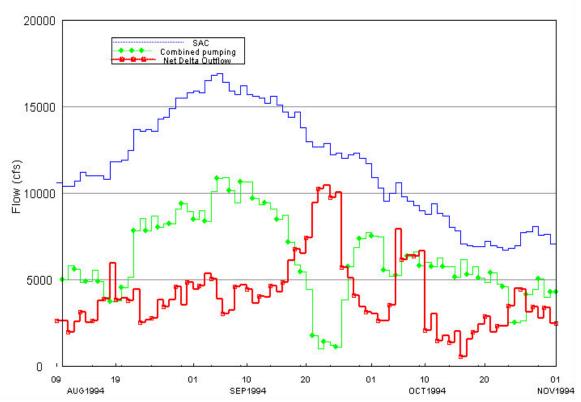


Figure 10.5: Sacramento River Flow, Combined Exports, and Net Delta Outflow in Fall 1994.

The incident has already been the subject of some retrospective water cost analysis by operators, though not with the use of optimization. They compared the abrupt historical reduction in pumping to that of a milder "adjusted" reduction begun earlier that was also adequate to achieve compliance. Figure 10.6 compares the actual and adjusted pumping strategies indirectly by means of the resulting Delta Outflow Index (DOI), a preliminary estimate of Net Delta Outflow (NDO). Sacramento inflow was held at historical levels. A comparison of the cumulative water cost of the two pumping strategies is shown in Figure 10.7. More detailed input and output from this study are not available.

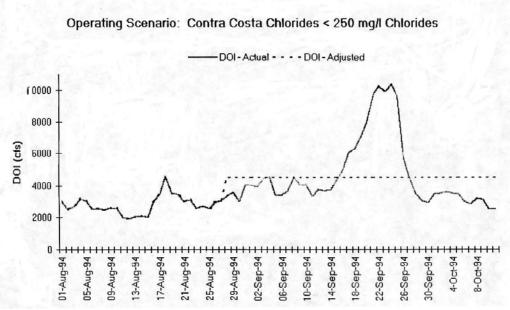


Figure 10.6: Historical and Alternative Delta Outflow Index for DWR O&M Fall 1994 Water Quality Study.

(taken from Creel, 2000)

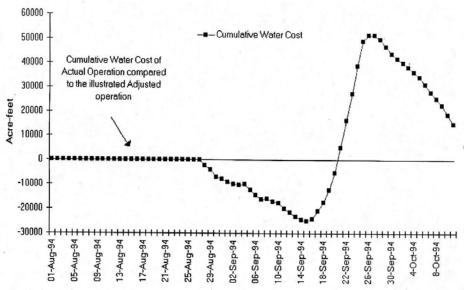


Figure 10.7: Comparison of the Cumulative Water Cost of Historical Operations to the Adjusted Operation for Fall 1994.

(taken from Creel, 2000)

Operators concluded that:

- □ the gentle method achieves compliance with a slightly lower water cost;
- □ the historical pulse achieved lower salinity with only a modest additional water cost.

The conclusion leaves open the question of the "cheapest" method, because the historical, pulse-like pump reduction had a higher cumulative water cost, but also overcompensated for salinity. Evidently, the historical pump reduction could been smaller or shorter (or timed differently).

Ateljevich (2001) extends the investigation to consider time-varying Sacramento flow and exports, using parameterized controls that are allowed to change every few days. The dynamics of the Delta were simulated using DSM2. The Sacramento River inflow and exports were treated as control variables, with other boundary conditions taken from the historical (IEP) data or DWR Delta island consumptive use (DICU) estimates. The "single step control" of 1,200 hours (50 days) was compared to finer grain controls of 50-hours and 100-hours. The "salinity regulation" is an daily average EC of 1,200 µmhos/cm on Old River at Holland Tract. This is a rough surrogate for the 250 mg/l chloride standard for the Contra Costa Canal intake at Rock Slough. Average pumping was fixed, although the pumping schedule was allowed to vary as a control variable. The experiment was repeated with average combined export (Central Valley Project and State Water Project) rates of 5,000 cfs and 7,000 cfs. The optimization was carried out using a successive linear programming algorithm based on work by Zhang et al. (1985)². Boundary EC at Martinez was estimated using the methods of Ateljevich (2001). Minimum outflow and E/I ratio regulations were ignored in order to study salinity control in isolation.

Figure 10.8 shows salinity and control solutions for the two levels of fixed pumping. Similar solutions were obtained from a variety of starting points. The optimization algorithm occasionally stalls, but when it does not, the minimum that it reaches is the same for all the starting points.

Just as in the simplified problem described in the previous section, the salinity trajectory lies very close to the regulation. Again, this was not a goal of the optimization, but is a reasonable result. The closeness of compliance is remarkable in light of how coarse the four-day controls are compared to the daily regulation. It is not well known in the Delta modeling community that such control over daily averaged salinity is possible, even in the sterile context of deterministic modeling.

¹ Two and four tide cycles, respectively. Due to limited computational power, the 50-hour controls contain some 100-hour periods later in the optimization period.

² See Ateljevich (2001) for more discussion of algorithms. This one is simple to program and is somewhat compatible with the linear programming engine of CALSIM. Somewhat more robust algorithms are available, although they tend to stall due to their use of finite differences.

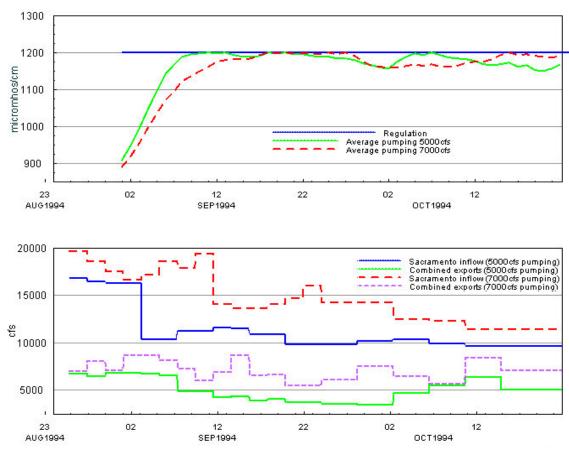


Figure 10.8: Salinity Results and Flows for Minimum Water Cost Optimization for two Cases: Average Pumping of 5,000 cfs and Average Pumping of 7,000 cfs.

The control solutions in Figure 10.8 are typical in that a decelerating pulse (high Sacramento flow and low combined exports) and relaxation is used to approach the salinity regulation at a tangent. The approach is gentler in the 7,000 cfs average pumping case than it is in the 5,000 cfs average pumping case. After the regulation is achieved, fairly steady values are observed. The spring-neap cycle and the accompanying filling and draining of the Delta have only a minor influence on the shape of the controls. One apparent influence of the spring-neap cycle or Delta filling and draining is the driving down of salinity before October 2 in order not to exceed the regulation as water levels in the Delta start to increase with the new spring tide (the tide cycle is shown in Figure 10.9).

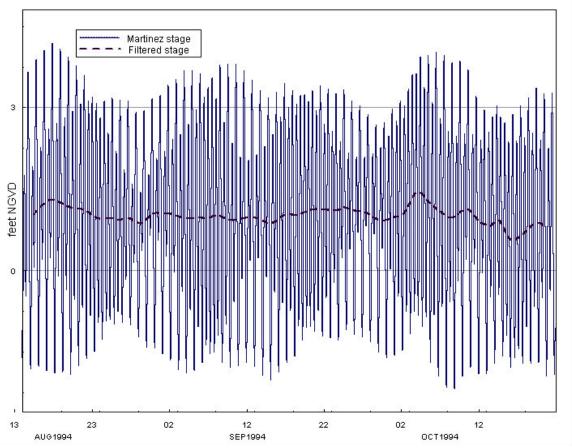


Figure 10.9: Martinez Stage and Tidal Average Using Godin Filter for Fall 1994.

Sacramento flow is more dynamic than exports. This is counter to conventional wisdom – pumping curtailment is usually treated as the main control action, on the (apparently questionable) assumption that control by means of Sacramento inflow is more expensive. Sacramento inflow tends to fluctuate in concert with exports: high flows are matched to low exports and vice versa (this is more evident in the 7,000 cfs average pumping case). This means that the E/I ratio fluctuates over a considerable range over the optimization period. In a production optimization in the Delta, the E/I ratio would have to be fixed as an additional constraint because it does not arise naturally from water quality control.

Finally, Ateljevich (2001) performs a sensitivity analysis on the time steps for inflow and pumping controls. Only controls coarser than one per day are possible for computational reasons. Given this constraint, it is interesting to question whether the solution is stable for various length control periods that are in the neighborhood of several days long and if so, at what level of detail the benefits tend to taper off. In the experiment presented here (from the 5,000 cfs average pumping case), the water cost from the 50-hour controls is only about 1% better than that resulting from the 100-hour controls and the shape of the control trajectory is similar. The incremental benefits of using a 50-hour time-step instead of a 100-hour time-step are meager compared to the 10% reduction in water costs obtained by using a 100-hour time-step instead of a single decision. The time-step sensitivity was not tested for the 7,000 cfs pumping case. However because the 7,000 cfs solution varies more over time, the difference between the

coarser and finer time steps might be larger than the 1% difference in the results for the 5,000 cfs pumping case.

10.5 Extensions

The optimization methods of Ateljevich (2001) have applications in both operational and planning models. In an operational context, optimization is able to provide an efficient reference solution for water cost. This reference solution can either be used as part of an operator decision based on multiple objectives, or it can be used as the "nominal trajectory" in a stochastic optimal control problem where the influence of uncertainty in the tide, inflows, and model are taken into account. For now, the main interest in optimal solutions will be heuristic. An example given above is the discovery that the control of salinity by means of extra Sacramento River flow may be more efficient than was previously thought.

The extension to planning models requires that the planning model CALSIM treat salinity control as a multiple-step procedure. This is a significant extension to CALSIM, and would require surrogate models (ANNs) that not only estimate salinity well at the "bottom line", but also correctly estimate the sensitivity of salinity to individual parts of the recent flow history. This is a more stringent fitting standard than the one currently in use.

The application of flow simulations to optimized CALSIM runs would also be facilitated by incorporating the formal theory of surrogates recently forwarded by Booker et al. (1998). Their methods formally tackle the idea of using a surrogate (ANNs are specifically addressed) to approximate a complicated component (DSM2) in an optimization procedure. The theory, while more expensive than a pure CALSIM-ANN implementation, is able to achieve convergence between the optimizer procedure (CALSIM) and the expensive model (DSM2) while using the surrogate for the overwhelming majority of the work.

10.6 References

- Ateljevich, E. (2001). *Optimal Control of Salinity in the Sacramento-San Joaquin Delta*. Ph.D. Dissertation. University of California, Berkeley.
- Booker, A.J., J.E. Dennis, Jr., P.D. Frank, D.B. Serafini, V. Torczon, and M.W. Trosset. (1998). A Rigorous Framework for Parameterization of Expensive Functions by Surrogates. NASA CR-1998-208735, ICASE Report TR-98-47.
- Creel, C. (2000). Personal communication concerning DWR Operations and Maintenance Delta compliance issues.
- Finch, R., and S. Kao. (1992). "Coarse-Grain Parallel Computing Using ISIS Tool Kit." *Journal of Computing in Civil Engineering*. 6 (2). pp. 233-44.
- Roache, P.J. (1992). "A Flux-Based Modified Method of Characteristic." *International Journal of Numerical Methods in Fluids.* 15.

Zhang, J., N. Kim, and L. Lasdon. (1985). "An Improved Successive Linear Programming Algorithm." *Management Science.* 31 (1).